arXiv:1606.05407v1 [stat.ME] 17 Jun 2016

Pyramid Quantile Regression

T. Rodrigues*†‡, J.-L. Dortet-Bernadet§ and Y. Fan*

Abstract

Quantile regression models provide a wide picture of the conditional distributions of the response variable by capturing the effect of the covariates at different quantile levels. In most applications, the parametric form of those conditional distributions is unknown and varies across the covariate space, so fitting the given quantile levels simultaneously without relying on parametric assumptions is crucial. In this work we propose a Bayesian model for simultaneous linear quantile regression. More specifically, we propose to model the conditional distributions by using random probability measures known as quantile pyramids. Unlike many existing approaches, our framework allows us to specify meaningful priors on the conditional distributions, whilst retaining the flexibility afforded by the nonparametric error distribution formulation. Simulation studies demonstrate the flexibility of the proposed approach in estimating diverse scenarios, generally outperforming other competitive methods. The method is particularly promising for modelling the extremal quantiles. Applications to linear splines and extreme value analysis are also explored through real data examples.

*School of Mathematics and Statistics, University of New South Wales, Sydney 2052 Australia.

†CAPES Foundation, Ministry of Education of Brazil, Brasília - DF 70040-020, Brazil

‡Communicating Author: t.rodrigues@unsw.edu.au

§Institut de Recherche Mathématique Avancée, UMR 7501 CNRS, Université de Strasbourg, Strasbourg, France.

Keywords: Bayesian quantile pyramid; Simultaneous quantile regression; Linear quantile regression; Extremal quantile regression.

1 Introduction

Since the seminal work by Koenker and Bassett (1978), linear quantile regression has been recognized in recent years as a robust statistical procedure that offers a powerful and compelling alternative to ordinary linear mean regression. It has been successfully applied to a diverse range of fields whenever interest lies in the non-central parts of the response distribution, often found in the environmental sciences, medicine, engineering and economics. Let τ, $0 < \tau < 1$, be a probability value and let $\mathcal{X}$ be a bounded subspace of $\mathbb{R}^P$, for an integer $P \geq 1$. The linear τ-th quantile regression model specifies the conditional distribution of a real response variable Y given the value $X = \mathbf{x}$ of a P dimensional vector of covariates

$$Y|\mathbf{x} \quad \sim \quad \beta_\tau^0 + \mathbf{x}'\beta_\tau + \epsilon, \tag{1}$$

for some unkown coefficients $\beta_\tau^0 \in \mathbb{R}$ and $\beta_\tau \in \mathbb{R}^P$, and for a noise variable ϵ whose τ-th conditional quantile is 0, *i.e.* $Q_\tau(\epsilon|\mathbf{x}) \equiv \inf\{a : P(\epsilon \leq a|X = \mathbf{x}) \geq \tau\} = 0$ or $\mathbb{P}(\epsilon \leq 0|X = \mathbf{x}) = \tau$. Equivalently, we can write the τ-th quantile of the conditional distribution of Y given $X = \mathbf{x}$ as $Q_\tau(Y|\mathbf{x}) = \beta_\tau^0 + \mathbf{x}'\beta_\tau$.

Let $(y_i, \mathbf{x}_i)_{i=1,\ldots,N}$ be N observed values of (Y, X). Frequentist inference on the linear quantile regression model typically leaves the noise distribution unspecified and the estimation of $(\beta_\tau^0, \beta_\tau)$ is carried out by solving the minimization problem,

$$(\hat{\beta}_\tau^0, \hat{\beta}_\tau) = \underset{(\beta^0,\beta)}{\arg\min} \sum_{i=1}^{N} \rho_\tau(y_i - \beta^0 - \mathbf{x}_i'\beta)\,,$$

where the so-called "check function" $\rho_\tau(.)$ is given by $\rho_\tau(\epsilon) = \tau\epsilon$ if $\epsilon \geq 0$ and $\rho_\tau(\epsilon) =$

$(\tau - 1)\epsilon$ otherwise (see Koenker and Bassett 1978). Inference is usually based on asymptotic arguments, see Koenker (2005) for more details and properties of this approach. Bayesian treatment of quantile regression is more challenging, since a specification of a likelihood can be problematic. In recent years, the asymmetric Laplace error model has emerged as a popular tool for Bayesian inference (Yu and Moyeed 2001), largely due to its flexibility and simplicity, and the fact that the corresponding maximum likelihood estimate is the solution of the minimization problem above. It was shown in Sriram et al. (2013) that, under mild conditions, the asymmetric Laplace can produce a Bayesian consistent posterior inference for the case of linear quantiles. However, in applications to real data, we do not really expect the distribution of the underlying data to follow an asymmetric Laplace distribution. Empirically, several authors have demonstrated that the asymmetric Laplace model does not have good coverage probabilities, see for example Reich et al. (2008). Other authors have tried to model the error distribution flexibly with nonparametric distributions, constraining the τ-th quantile of the error distribution to be zero. See *e.g.* Kottas and Gelfand (2001), Hanson and Johnson (2002), Kottas and Krnjajić (2009) or Reich et al. (2008) who propose the use of various nonparametric distributions including infinite mixture of Gaussians, Dirichlet process mixtures and mixture of Pólya trees.

In many applications, quantile estimates at several different quantile levels are needed to provide a precise description of the conditional distribution. A well known problem with separately fitted quantile regression planes is that they can cross, violating the definition of quantiles. A possible solution is to use a second stage adjustment to the initial fits, see for example Hall et al. (1999), Dette and Volgushev (2008) and Chernozhukov et al. (2009) in the frequentist setting, or more recently Rodrigues and Fan (2016) in a Bayesian setting. Another possible solution is a joint estimation of multiple quantiles. This has been advocated by several authors, as it leads naturally to a greater borrowing of information across quantiles and a higher global efficiency for all quantiles of interest.

Under this paradigm, Reich et al. (2011) proposed a model using Bernstein basis polynomials for spatial quantile regression. Tokdar and Kadane (2012) and Yang and Tokdar (2016) treat the regression coefficients as a function of τ, using smooth monotone curves to model them under a Gaussian process prior. One of the common issues facing the more general modelling approach is that the likelihood is not available in analytic form, leading to the necessity to numerically approximate the likelihood values for each data observation. This has the computational disadvantage of an added source of error from the numerical approximation, and sometimes prohibitive computational cost when the dataset is large, see for example Reich et al. (2011). A closed-form likelihood approach is proposed by Reich and Smith (2013), who extended the location scale model of He (1997) to more flexibly model the quantile process. More recently, Fang et al. (2015) proposed to use a linearly interpolated approximate likelihood derived from the quantiles, where the peudo-likelihood is available in analytical form, which approaches the true likelihood with increasing number of quantiles.

In this paper, our idea is to parameterise quantile regression models directly with quantiles. To do this, we make use of the quantile pyramids described in Hjort and Walker (2009) for constructing random probability measures. There are several advantages to our modelling construct. Firstly, working with the quantiles provides an easily interpretable model setup which allows us to place sensible priors on the conditional distributions. Secondly, our likelihood is available in closed form. Finally, the framework allows us to fit a single quantile or as many specific levels of quantiles as we want, making this the only method to do so in the literature as far as we are aware of.

The rest of the article is organised as follows. In Section 2 we recall the basic construction of the quantile pyramids studied in Hjort and Walker (2009). The proposed pyramid quantile regression (PQR) modelling is detailed in Section 3, including its theoretical properties and an estimation procedure. Extensive simulation studies are carried out in Section 4, where the proposed method is compared to the best alternative approaches. In

Section 5, real examples illustrate PQR application to quantile linear splines and extreme quantile modelling. The final section presents concluding discussions.

2 Quantile pyramids for random distributions

Quantile pyramids was introduced by Hjort and Walker (2009) as a method to define a random probability measure for nonparametric Bayesian inference. Contrary to the better known Pólya trees (Ferguson 1974, Lavine 1992, Lavine 1994) that consider random probability masses and fixed partitions, Hjort and Walker (2009) propose the use of random quantiles with fixed probabilities.

The pyramid quantile process that defines a random probability measure on $[0, 1]$ is constructed as follows. Let $Q(\tau)$ be the associated random quantile function, with $Q(0) = 0$ and $Q(1) = 1$. At level $m = 1$ of the construction the median $Q(1/2)$ is randomly generated over $(0, 1)$ according to a given distribution. At level $m = 2$ of the construction the quartile $Q(1/4)$ is sampled on the interval $(0, Q(1/2))$ and $Q(3/4)$ is sampled over $(Q(1/2), 1)$. The process is continued at the following levels m, where the quantiles $Q(j/2^m), j = 1, 3, \ldots, 2^m - 1$, are generated conditionally on the quantiles previously sampled. Figure 1(a) demonstrates one sample drawn from this quantile pyramid process for $m = 1, 2, 3$, where the value of $Q(j/2^m)$ is indicated on the x-axis, and Figure 1(b) shows the intervals from which successive quantiles at different levels were sampled.

Specifically, quantiles at level m are generated after those at level $m - 1$ according to

$$Q\left(j/2^m\right) = Q\left((j-1)/2^m\right)\left(1 - V_{mj}\right) + Q\left((j+1)/2^m\right)V_{mj}\,, \tag{2}$$

where $Q\left(j/2^m\right)$ is the new quantile defined at level m and where $Q\left((j-1)/2^m\right)$ and $Q\left((j+1)/2^m\right)$ are its closest ancestors. The independent variable at work at each level m, V_{mj}, is a random variable on the unit interval. A natural choice is to use V_{mj}'s that are Beta distributed, see Hjort and Walker (2009) for other possibilities. As m tends to infinity

the random quantile $Q(\tau)$ is defined for all τ in $(0, 1)$. Notably, the behaviour of this quantile pyramid process depends on these variables. For instance, if at each level we impose that $\mathbb{E}(V_{mj}) = 0.5$ then we have $\mathbb{E}(Q(\tau)) = \tau$ for all τ in $(0, 1)$ and the quantile process is centered at the uniform quantile function. Theoretical results that concern $Q(\tau)$ can be found in Hjort and Walker (2009). They describe for example relatively mild conditions involving decreasing variances of V_{mj} for growing m that ensure *a.s.* the existence of an absolutely continuous $Q(\tau)$.

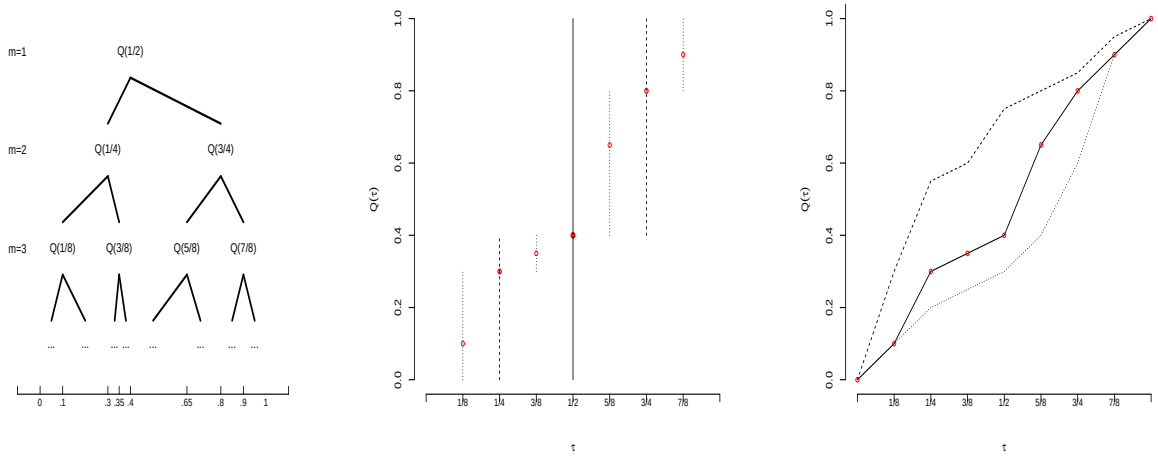

(a) (b) (c)

Figure 1: The quantile pyramid generating process. Figure (a) shows the binary tree for one sample drawn from this quantile pyramid process, x-axis indicates the value of $Q(j/2^m)$. In figure (b), the lines indicate the intervals from which the quantile values were sampled. Figure (c) shows the interpolated quantile function, the different curves correspond to different samples of the quantile function.

In practice, to allow a Bayesian inference on the random distribution, the process is stopped at a finite level M and a linear interpolation on the set of quantiles $Q(j/2^M)$, $j = 0, 1, ..., 2^M$, completes the process. Figure 1(c) demonstrates three random samples of the piecewise linear quantile functions obtained from the described procedure. The density function corresponding to this linearly interpolated quantile function is piecewise constant, so there is a well defined likelihood function for this random type histogram

model. Due to the tree nature of the quantile process, the simultaneous density of the $2^M - 1$ quantiles can be written as

$$\begin{aligned} &\pi\left(Q\left(\tfrac{1}{2}\right), Q\left(\tfrac{1}{4}\right), Q\left(\tfrac{3}{4}\right), \ldots, Q\left(\tfrac{2^{M-1}}{2^M}\right)\right) \\ &\qquad = \prod_{m=1}^{M} \left\{ \prod_{j=1,3,\ldots,2^{m-1}} \pi_{mj}\left(Q\left(\tfrac{j}{2^m}\right) \mid Q\left(\tfrac{j-1}{2^m}\right), Q\left(\tfrac{j+1}{2^m}\right)\right) \right\}, \end{aligned} \tag{3}$$

where the densities π_{mj} can be derived from Equation 2, based on the density of V_{mj}, through a simple transform of variables.

3 Regression modelling with quantile pyramids

Here we introduce the use of quantile pyramids in the linear regression setting. We consider the general case when several conditional quantiles are of interest, say $Q_\tau(Y|\mathbf{x})$ at quantile levels $\tau = \tau_1, \tau_2, \ldots, \tau_T$ with $\tau_1 < \tau_2 < \ldots < \tau_T$. The covariate $\mathbf{x} = (x_1, ..., x_P)$ belongs to a given bounded subset $\mathcal{X}$ of $\mathbb{R}^P$. In practice $\mathcal{X}$ can be taken as the convex hull of the N observed data points $\mathbf{x}_i$, $i = 1, ..., N$.

3.1 Model formulation

The starting point for the model formulation is the simple fact that a hyperplane in $\mathbb{R}^{P+1}$ is determined by the values of $P+1$ of its points. Let $\mathbf{x}^0, \mathbf{x}^1, \ldots, \mathbf{x}^P$ denote any $P+1$ locations with corresponding τth conditional quantile denoted by $Q^p_\tau, p = 0, \ldots, P$. Without loss of generality let $\mathbf{x}^0 = (0, ..., 0)$, $\mathbf{x}^1 = (1, 0, ..., 0)$, $\mathbf{x}^2 = (0, 1, 0, ..., 0)$, ..., $\mathbf{x}^P = (0, ..., 0, 1)$. The linear quantile regression model for the τth conditional quantile $Q_\tau(Y|\mathbf{x})$ can be described by the hyperplane passing through these $P + 1$ points

$$\begin{aligned} Q_\tau(Y|\mathbf{x}) &= Q^0_\tau + \textstyle\sum_{p=1}^{P}(Q^p_\tau - Q^0_\tau)x_p \\ &\equiv \beta_0(\tau) + \textstyle\sum_{p=1}^{P} \beta_p(\tau)x_p \,, \end{aligned} \tag{4}$$

where $\beta_0(\tau)$ and $\beta_p(\tau)$ denote the regression coefficients at $\tau = \tau_1, \tau_2, \ldots, \tau_T$. For other other choices of locations $\mathbf{x}^0, ..., \mathbf{x}^P$, Equation 4 which is simply the equation of a plane passing through these points has to be modified. In short, the proposed model for simultaneous linear quantile regression uses $P + 1$ independent finite pyramid quantile processes for the quantile functions Q^p_τ. Before proceeding to describe the likelihood, we first present some extensions of these processes that are important in the quantile regression context.

3.2 Oblique quantile pyramid

The quantile pyramid described in Section 2 uses a dyadic partitioning of the probability interval $[0, 1]$. In this setting, the induced quantile levels are all fixed and equally spaced. However, in practice, we may be interested in quantiles at specific levels τ over an irregular grid.

In these circumstances, the quantile level of a child node of the pyramid tree is usually no longer located in the middle point of the quantile levels of its closest ancestors. We call this general setting oblique quantile pyramid, as opposed to the regular pyramid previously described. To keep the process centred on the Uniform distribution, we now choose $E(V_{mj})$ to reflect this unequal split using the relative distance from the child quantile level τ_{mj} to its closest ancestors,

$$E(V_{mj}) = \frac{\tau_{mj} - \tau^L_{mj}}{\tau^R_{mj} - \tau^L_{mj}}\,, \tag{5}$$

where τ^L_{mj} and τ^R_{mj} denote its left and right nearest ancestors' quantile levels, respectively. From Equations 2 and 5, it is easy to see that $E(Q(\tau)) = \tau$, i.e. under this construction the oblique quantile pyramid is also centred on the Uniform distribution.

The oblique pyramid is constructed via the following procedure. For a sequence of quantiles $Q(\tau_t), t = 1, \ldots, T$, the first level of the pyramid at $m = 1$ generates the quantile whose level is halfway into the set of given quantile levels, we will call it the middle

quantile level (not to be confounded with the classic median quantile). If T is odd, this is $Q(\tau_{[T/2]+1})$, and given that $V_{11} \sim Beta(\alpha_{11}, \beta_{11})$, we set α_{11}, β_{11} such that $E(V_{11}) = \tau_{[T/2]+1}$, as $\tau_{11}^L = 0$ and $\tau_{11}^R = 1$ per construction. For the next level $m = 2$, we proceed by getting the middle quantile levels from the left and right of $Q(\tau_{[T/2]+1})$ to be the next nodes, and choose the corresponding $\alpha's, \beta's$ to satisfy Equation 5. The process is then continued until all quantiles in the sequence $Q(\tau_t), t = 1, \ldots, T$, have been specified. For identification purposes, if we have an even number of quantile levels, we define the middle value to be the smallest of the two middle quantile levels.

In addition, we choose to have the parameters α_{mj}, β_{mj} increasing with the pyramid level m, which reduces the prior variance for growing m. Throughout this paper, we choose $\alpha_{mj} = 2m$ and $\beta_{mj} = \alpha_{mj} * (1 - E(V_{mj})/E(V_{mj})$, if $E(V_{mj}) < 0.5$, where $E(V_{mj})$ is calculated using Equation 5. Otherwise, considering the symmetric nature of the Beta distribution, if $E(V_{mj}) \geq 0.5$, we take $\beta_{mj} = 2m$ and $\alpha_{mj} = \beta_{mj} * E(V_{mj})/(1 - E(V_{mj}))$. From our experience, this prior is not very informative and gives a good mixing in Markov chain Monte Carlo (MCMC) posterior simulations.

3.3 Centering the prior

Using random quantile functions $Q_\tau^p, p = 0, 1, \ldots, P$ for the linear model in Equation 4 defines a prior over the quantile planes. This prior should reflect the prior knowledge with respect to the response Y. The pyramid quantile building process described in Section 3.2 is centered on the Uniform distribution on $[0, 1]$. Let $Q_\tau^{p,unif}, p = 0, 1, \ldots, P$, be $P + 1$ independent replications of this process. In order to use the pyramid quantiles in Equation 4, for data Y arising from the reals, we can centre each Q_τ^p process on the quantile function of a normal distribution $\mathcal{N}(\mu^p, (\sigma^p)^2), p = 0, ..., P$, via a simple transformation suggested in Hjort and Walker (2009),

$$Q_\tau^p = \mu^p + \sigma^p \Phi^{-1}(Q_\tau^{p,unif}), \tag{6}$$

where Φ^{-1} denotes the quantile function of the standard normal distribution, for some mean parameters μ^p and standard deviation parameters σ^p. In this case, for each τ in $(0,1)$, the median of the random quantile Q^p_τ is the τth quantile of a normal distribution $\mathcal{N}(\mu^p, (\sigma^p)^2)$. More generally one can centre the prior on different distributions other than the Normal, depending on the specific prior knowledge available for the particular application at hand, by setting $Q^p_\tau = Q^*(Q^{p,unif}_\tau)$ for some arbitrary quantile function Q^*. Centring the prior on appropriate distributions can be particularly useful for estimating extreme quantiles, as data is scarce at the tails and the pyramid prior is more informative in the tails. However, it is our experience that, for non-extreme quantiles, results are not very sensitive to the default choice of the Normal distribution. For the clarity of exposition, we use a prior of the form of Equation 6 for the pivotal quantile pyramids $Q^p_\tau, p = 0, \ldots, P$ to describe our methodology.

In the finite quantile pyramid context a random density for Q^p_τ can be derived, which is piecewise scaled Normal distribution between the quantiles $Q^p_{\tau_1}, \ldots, Q^p_{\tau_T}$. This density is obtained by using a simple change of variable on the piecewise constant density function corresponding to $Q^{p,unif}_\tau$. Figure 2 illustrates some samples of this quantile process, highlighting the piecewise Normal density feature. The examples were simulated from a pyramid process centred on the standard Normal distribution, with $M = 3$, $\tau = 0.125, 0.25, \ldots, 0.875$ and $V_{mj} \sim Beta(a, a)$, for $a = 1$ and 10.

3.4 Likelihood and posterior

Equation 4 gives the desired quantiles of the conditional distribution of Y given $X = \mathbf{x}$, with cdf $F(y|\mathbf{x})$. When priors of the form 6 are used, we need to define the likelihood function. The chosen option here is to consider that the density $f(y|\mathbf{x})$ of the conditional

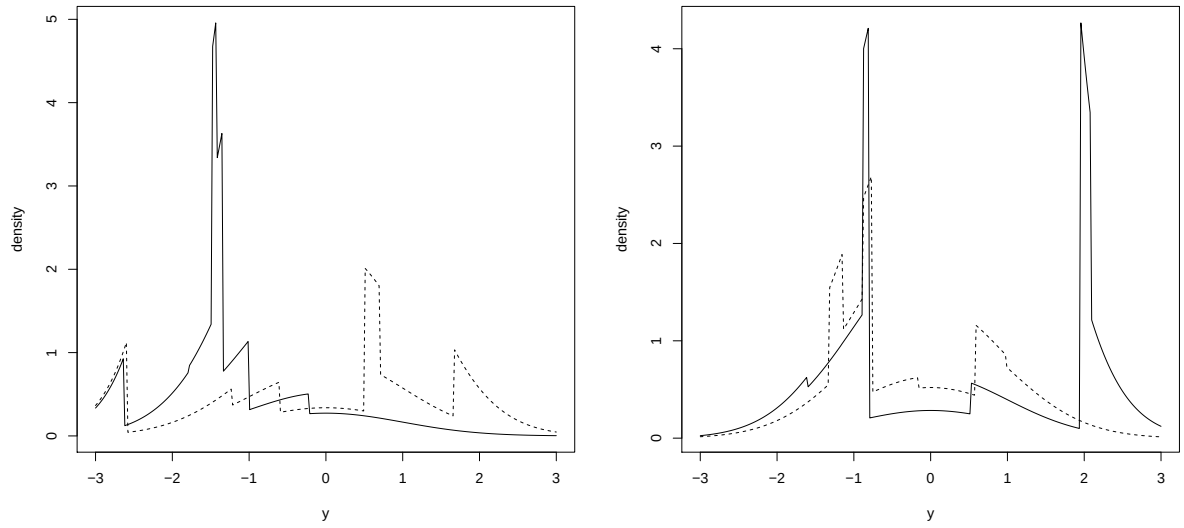

(a) (b)

Figure 2: Examples of simulated densities from a finite pyramid process centred on the standard Normal distribution, with $M = 3$ ($\tau = 0.125, 0.25, \ldots, 0.875$) and $V_{mj} \sim Beta(a, a)$, for $a = 1$ (left) and $a = 10$ (right).

distribution is piecewise Normal

$$f(y|\mathbf{x}) = \sum_{t=1}^{T} (\tau_t - \tau_{t-1}) \frac{\phi(y; \mu_{\mathbf{x}}, \sigma_{\mathbf{x}}^2)}{\Phi\left(\frac{Q_{\tau_t}(Y|\mathbf{x}) - \mu_{\mathbf{x}}}{\sigma_{\mathbf{x}}}\right) - \Phi\left(\frac{Q_{\tau_{t-1}}(Y|\mathbf{x}) - \mu_{\mathbf{x}}}{\sigma_{\mathbf{x}}}\right)} I_{(Q_{\tau_t}(Y|\mathbf{x}), Q_{\tau_{t-1}}(Y|\mathbf{x}))}(y), \qquad (7)$$

where $I_{(q_1,q_2]}(y)$ is 1 if $y \in (q_1, q_2]$ and zero otherwise, where $\phi(\cdot; \mu, \sigma^2)$ denotes the density function of a Normal distribution $\mathcal{N}(\mu, \sigma^2)$ and where the parameters $\mu_{\mathbf{x}}$ and $\sigma_{\mathbf{x}}$ change linearly in $\mathbf{x}$

$$\mu_{\mathbf{x}} = \left(1 - \sum_{p=1}^{P} x_p\right) \mu^0 + \sum_{p=1}^{P} x_p \mu^p \,, \qquad \sigma_{\mathbf{x}} = \left(1 - \sum_{p=1}^{P} x_p\right) \sigma^0 + \sum_{p=1}^{P} x_p \sigma^p \,. \qquad (8)$$

This formulation implies that the priors on all conditional distributions are centred on the Normal distribution and Equation 8 specifies that the quantiles of these centring distributions change linearly in the covariates. Equation 7 can be obtained by extending the linearly interpolated likelihood corresponding to the finite quantile pyramid centred on

the uniform distribution as in Hjort and Walker (2009) and applying the relevant transformation of the Equation 6.

Note that an alternative approach is to work directly with the density of the conditional distribution $f(y|\mathbf{x}) = \frac{1}{Q'_{F(y_i|\mathbf{x})}(Y|\mathbf{x})}$, where $Q'_\tau(Y|\mathbf{x})$ denotes the derivative of $Q_\tau(Y|\mathbf{x})$ with respect to τ. This derivative can be deduced from Equation 4. Nevertheless the value $F(y|\mathbf{x})$ can only be approximated and the involved numerical cost to approximate $f(y_i|\mathbf{x}_i)$ for each observation $i = 1, ..., N$ can be prohibitive. Note also that the implied conditional density of this approach is not piecewise normal at each $\mathbf{x}$.

The posterior distribution for the finite number of quantile levels $\tau_1, \ldots, \tau_T$ can be obtained for the quantiles $\mathbf{Q}^p = \{Q^p_{\tau_1}, \ldots, Q^p_{\tau_T}\}, p = 0, \ldots, P$, and the associated parameters $\boldsymbol{\mu} = \{\mu^0, \ldots, \mu^P\}, \boldsymbol{\sigma} = \{\sigma^0, \ldots, \sigma^P\}$, via the usual Bayes theorem

$$\pi(\mathbf{Q}^0, \ldots, \mathbf{Q}^P, \boldsymbol{\mu}, \boldsymbol{\sigma}|y_1, \ldots, y_N) \propto \prod_{i=1}^{N} f(y_i|\mathbf{x}_i) \times \prod_{p=0}^{P} \pi(\mathbf{Q}^p|\mu^p, \sigma^p) \times \pi(\boldsymbol{\mu}) \times \pi(\boldsymbol{\sigma}), \quad (9)$$

where $f(y_i|\mathbf{x}_i)$ is given in Equation 7. The distributions $\pi(\boldsymbol{\mu})$ and $\pi(\boldsymbol{\sigma})$ are hyperpriors for the parameters of the Normal distributions. Throughout the paper these hyperpriors are set to $N(0, 20)$ and Gamma$(0.001, 0.001)$, for μ^p and σ^p respectively. In addition, using Equations 3 and 6, the pivotal pyramid prior distributions $\pi(\mathbf{Q}^p|\mu^p, \sigma^p)$, $p = 0, \ldots, P$, are

$$\begin{aligned}\pi(\mathbf{Q}^p|\mu^p, \sigma^p) &= \prod_{m,j} \pi_{mj} \left(Q^p_{\tau_{mj}} | Q^p_{\tau^L_{mj}}, Q^p_{\tau^R_{mj}} \right) \\ &= \prod_{m,j} \left\{ g\left(\frac{\Phi\left(\frac{Q^p_{\tau_{mj}} - \mu^p}{\sigma^p}\right) - \Phi\left(\frac{Q^p_{\tau^L_{mj}} - \mu^p}{\sigma^p}\right)}{\Phi\left(\frac{Q^p_{\tau^R_{mj}} - \mu^p}{\sigma^p}\right) - \Phi\left(\frac{Q^p_{\tau^L_{mj}} - \mu^p}{\sigma^p}\right)} \right) \times \frac{\phi\left(\frac{Q^p_{\tau_{mj}} - \mu^p}{\sigma^p}\right)}{\Phi\left(\frac{Q^p_{\tau^R_{mj}} - \mu^p}{\sigma^p}\right) - \Phi\left(\frac{Q^p_{\tau^L_{mj}} - \mu^p}{\sigma^p}\right)} \right\},\end{aligned}$$

where $g(\cdot)$ denotes the density of the V_{mj} variables, throughout the paper $V_{mj} \sim Beta(\alpha_{mj}, \beta_{mj})$, and ϕ the standard Normal density.

3.5 Non-crossing constraints

The the linear model proposed in this paper ensures that the simultaneously fitted quantile planes in Equation 4 do not cross on the convex hull of the $P + 1$ pivotal locations $\mathbf{x}^0, \mathbf{x}^1, \cdots, \mathbf{x}^P$. For the single covariate problem, choosing pivotal locations $\mathbf{x}^0$ and $\mathbf{x}^1$ to be the minimum and maximum value of $\mathcal{X}$ is sufficient to ensure non-crossing. However, for $P > 1$, some caution is needed with crossings. If $\mathcal{X}$ corresponds to the convex hull of the observed data points and the pivotal quantiles are placed at $P + 1$ well separated vertices, the non-crossing of the planes needs to be verified at any remaining vertices of the convex hull, say at some points denoted $\mathbf{x}^e$'s. Note that working with any convex sets larger than the minimum convex set enclosing the data will also ensure non-crossing, but convex sets that are too large puts unnecessary constraints on the regression model, forcing the regression planes to be parallel.

A naive option to ensure non-crossing is to check for crossing at the non-pivotal locations of the convex hull and discard the samples that produce crossing planes during the Metropolis-Hastings MCMC sampling procedure. However, for moderate numbers of covariates, this approach is very inefficient as the crossing will most likely be frequent. In fact, a better solution is to adjust the MCMC proposal distribution so that it proposes only in the non-crossing region. To accomplish that, we will adopt Uniform proposals $U(l_\tau^p, u_\tau^p)$ for Q_τ^p, and choose the lower and upper bounds (l_τ^p, u_τ^p) while ensuring that the corresponding quantiles at the non-pivotal locations do not cross, Fang et al. (2015) used a similar approach working with the entire dataset.

More specifically, for each extra location $\mathbf{x}^e$ on the vertices of the convex hull, the bounds can be easily found by solving $Q_\tau^e = Q_{\tau-1}^e$ and $Q_\tau^e = Q_{\tau+1}^e$ for Q_τ^p based on the hyperplane equation. For example, for the hyperplane described in Equation 4, when the

pth component of $\mathbf{x}^e$ is greater than zero, i.e. $x_p^e > 0$, we have

$$l_\tau^p|\mathbf{x}^e = Q_\tau^0 + \left(Q_{\tau-1}^e - Q_\tau^0 - \sum_{j \neq p} \left(Q_\tau^j - Q_\tau^0 \right) x_j^e \right) / x_p^e \,,$$
$$u_\tau^p|\mathbf{x}^e = Q_\tau^0 + \left(Q_{\tau+1}^e - Q_\tau^0 - \sum_{j \neq p} \left(Q_\tau^j - Q_\tau^0 \right) x_j^e \right) / x_p^e \,,$$

where $l_\tau^p|\mathbf{x}^e$ and $u_\tau^p|\mathbf{x}^e$ are Q_τ^p lower and upper bounds, respectively, based on the crossing restrictions at $\mathbf{x}^e$. Similarly, if $x_p^e < 0$, the above lower bound becomes the upper bound and vice-versa. Therefore, to take into account all non-pivotal vertices' constraints, we choose $l_\tau^p = \max\{l_\tau^p|\mathbf{x}^e\}$ and $u_\tau^p = \min\{u_\tau^p|\mathbf{x}^e\}$. In this way, non-crossing issues are easily handled. Note that the bounds for each extra location can be found at once through simple matrix operations, not being computationally very expensive.

3.6 Posterior consistency

Consistency properties of the proposed method can be derived from the properties of pyramid quantiles.

Consider the finite quantile pyramid construction for a single population introduced in Section 2. This process samples until level M a finite set of quantile values, say $\theta = \{Q_{\tau_1}, ..., Q_{\tau_T}\}$. The other quantiles being defined by linear interpolation, it defines a parametric model $f(., \theta)$ for the data. The discussion in Section 5.1 of Hjort and Walker (2009) tells that, as the number N of observations increases, the posterior distribution for θ will concentrate on decreasing neighbourhoods around the "least false parameter", θ_0, defined as the minimiser in θ of the Kullback-Leibler distance from true model f to parametric model $f(., \theta)$ (see Hjort 1992). The proof is simple and considers that, for a prior with nonzero mass at θ_0, the Bayes estimator converges to the maximum likelihood estimator, which in turn, asymptotically, minimises the KL distance. This holds regardless of the centring distribution of the quantile process (for more general proofs, see Hjort 1992

and Strawderman and Tsiatis 1996).

Consequently, assuming that the linear model assumption is true, and given that, at the pyramid locations, the quantile pyramid estimates Q_τ^p are consistent for the least false parameters, then the posterior quantile estimates from the regression model (Equation 9) are also consistent. Cube root asymptotics govern the behaviour of those estimators and more details about the limiting process can also be found in Moss (2015).

Furthermore, under some mild conditions on the pyramid prior, primarily related to mean and variance of V_{mj}, and for pyramid levels M increasing with sample size N, the least false distribution converges to the true density. The sequence of posterior distributions is then Hellinger consistent (Hjort and Walker 2009) and the differences between θ_0 and the true quantiles go to zero. Although consistency to the true quantiles can only be obtained with growing M, for fixed M, in which only consistency to the least false quantiles holds, very good empirical results are produced. This is demonstrated next in our simulation studies, which explore quantile processes with small pyramid levels M.

4 Simulated examples

In this section, small sample properties of the pyramid quantile regression estimator (PQR) will be investigated through simulation examples. Also, PQR will be compared with three other approaches: semiparametric regression model (BSquare) of Reich and Smith (2013), Gaussian process method (GPQR) of Yang and Tokdar (2016) and the frequentist constrained estimator (freqQR) of Bondell et al. (2010). The comparisons will be undertaken in terms of the empirical root mean squared error $RMSE(\tau) = \sqrt{1/s \sum_{s=1}^{S} [\beta(\tau) - \hat{\beta}_s(\tau)]^2}$ and 95% coverage probabilities, based on $S = 200$ data sets. Following Reich and Smith (2013), we use the simulation designs that are detailed below.

Design 1. $\beta_0(\tau) = \log[\tau/(1-\tau)]$, $\beta_1(\tau) = 2$;

Design 2. $\beta_0(\tau) = \text{sign}(0.5-\tau)\log(1-2|0.5-\tau|)$, $\beta_1(\tau) = 2\tau$;

Design 3. $\beta_0(\tau) = \Phi^{-1}(\tau)$, $\beta_1(\tau) = 2\min\{\tau - 0.5, 0\}$;

Design 4. $\beta_0(\tau) = 2\Phi^{-1}(\tau)$, $\beta_1(\tau) = 2\min\{\tau - 0.5, 0\}$, $\beta_2(\tau) = 2\tau$, $\beta_3 = 2$, $\beta_4 = 1$, $\beta_5 = 0$;

For each design, we simulated some observations y_i, $i = 1, ..., N$, from

$$y_i = \beta_0(u_i) + \sum_{j=1}^{P} x_{ij}\beta_j(u_i),$$

where the j-th covariate is $x_{ij} \overset{iid}{\sim} \text{Unif}(-1, 1)$ and $u_i \overset{iid}{\sim} \text{Unif}(0, 1)$. The simulated conditional densities at $x = -1$, $f(Y|x = -1)$, for designs 1 to 4 are illustrated in Figure 3.

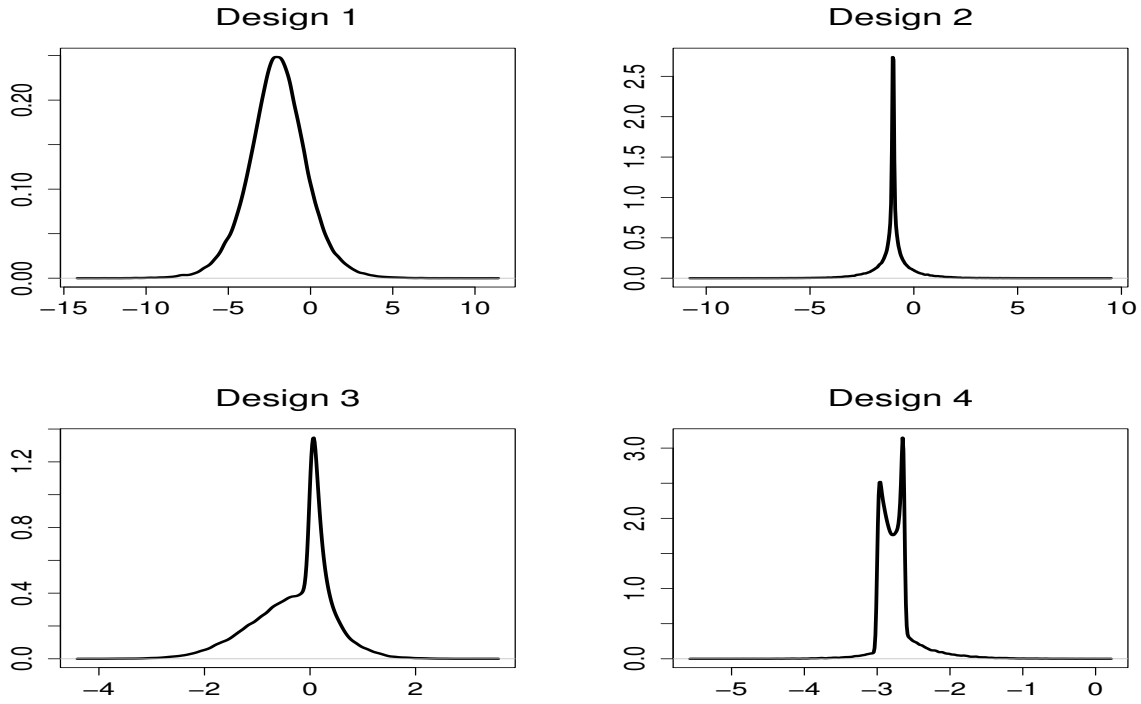

Figure 3: True conditional densities $f(Y|x = -1)$ for the simulation designs.

For univariate designs 1 to 3, we used the datasize $N = 100$ and estimated simultaneously the quantile regression lines at quantile levels $\tau = 0.01, 0.05, 0.10, ..., 0.95, 0.99$. PQR was fitted based on 110.000 MCMC draws and burn-in of 10.000. Furthermore, in order to improve MCMC mixing, the pyramid quantiles were reparametrised using the logarithm of the difference between adjacent quantile levels, i.e. $\{\log(Q^p(\tau_2) - Q^p(\tau_1)), \ldots,$

$\log\left(Q^p(\tau_T) - Q^p(\tau_{T-1})\right), \log\left(Q^p(\tau_1) + Q^p(\tau_T) + c\right)\}$, where a constant $c = 2|\min(Y_i)|$ was added to the last term to prevent a negative argument in the logarithm function.

BSquare estimator is implemented in BSquare package (Smith and Reich 2013) in R (R Core Team 2014), to fit this model we used the logistic base distribution with 4 basis functions. GPQR model is also available in R (qrjoint package by Tokdar 2015), and it was estimated from 50.000 MCMC samples, thinning every 10 samples and discarding the initial 20% of the samples as burn-in. Codes for Bondell et al. (2010) are available from first author's web page.

Figure 4 presents RMSE results for the univariate designs. Overall we can see that, for non-extreme quantile levels, all methods perform similarly, with BSquare having the best results for β_1 from Design 1 and PQR having the best results for β_1 from Design 3. Data from design 1 follows BSquare model assumptions, which certainly contributes to its better performance. Design 3 presents a more challenging quantile function, and the flexibility of the proposed approach is an advantage here.

For extreme quantiles, PQR clearly outperforms the other methods for most cases. Once again the flexibility of the proposed approach contributes to this achievement, as well as the reasonable choice of the quantile process centring distribution. Note that, although the simulated designs are not from a Normal distribution (e.g. see Figure 3), yet this is a reasonable centring choice here. The meaningfulness of quantile parameters in PQR is a great feature of the proposed model, as prior information are easily interpreted and incorporated. As shown here, a rough idea of the true distribution can contribute to improve the estimation of extreme quantiles.

Figure 5 shows 95% coverage probabilities at $\tau = 0.01, 0.05, 0.1, \ldots, 0.95, 0.99$ for the univariate designs. However, freqQR confidence intervals for the parameters at $\tau = 0.01$ and 0.99 are not available for this sample size, so freqQR results in Figure 5 are truncated at $\tau = 0.05$ and 0.95, and highlighted by diamond endpoints.

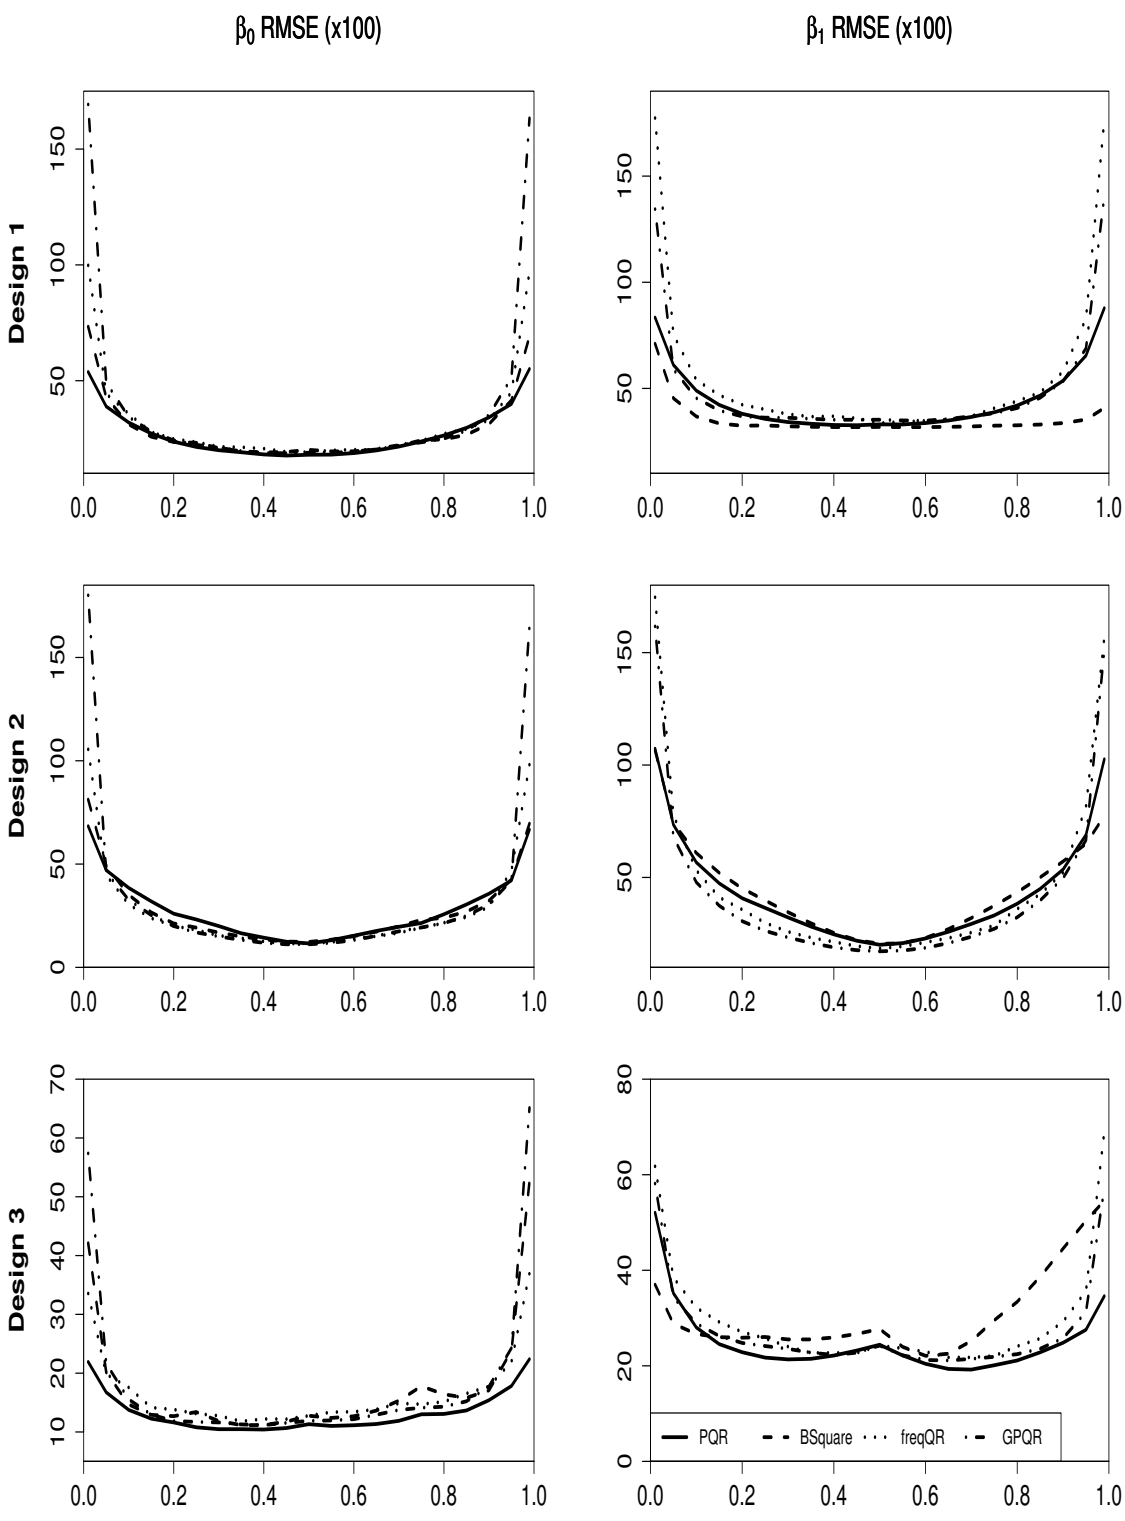

Figure 4: RMSE ($\times 100$) for β_0 (left) and β_1 (right) at $\tau = 0.01, 0.05, 0.1, \ldots, 0.95, 0.99$.

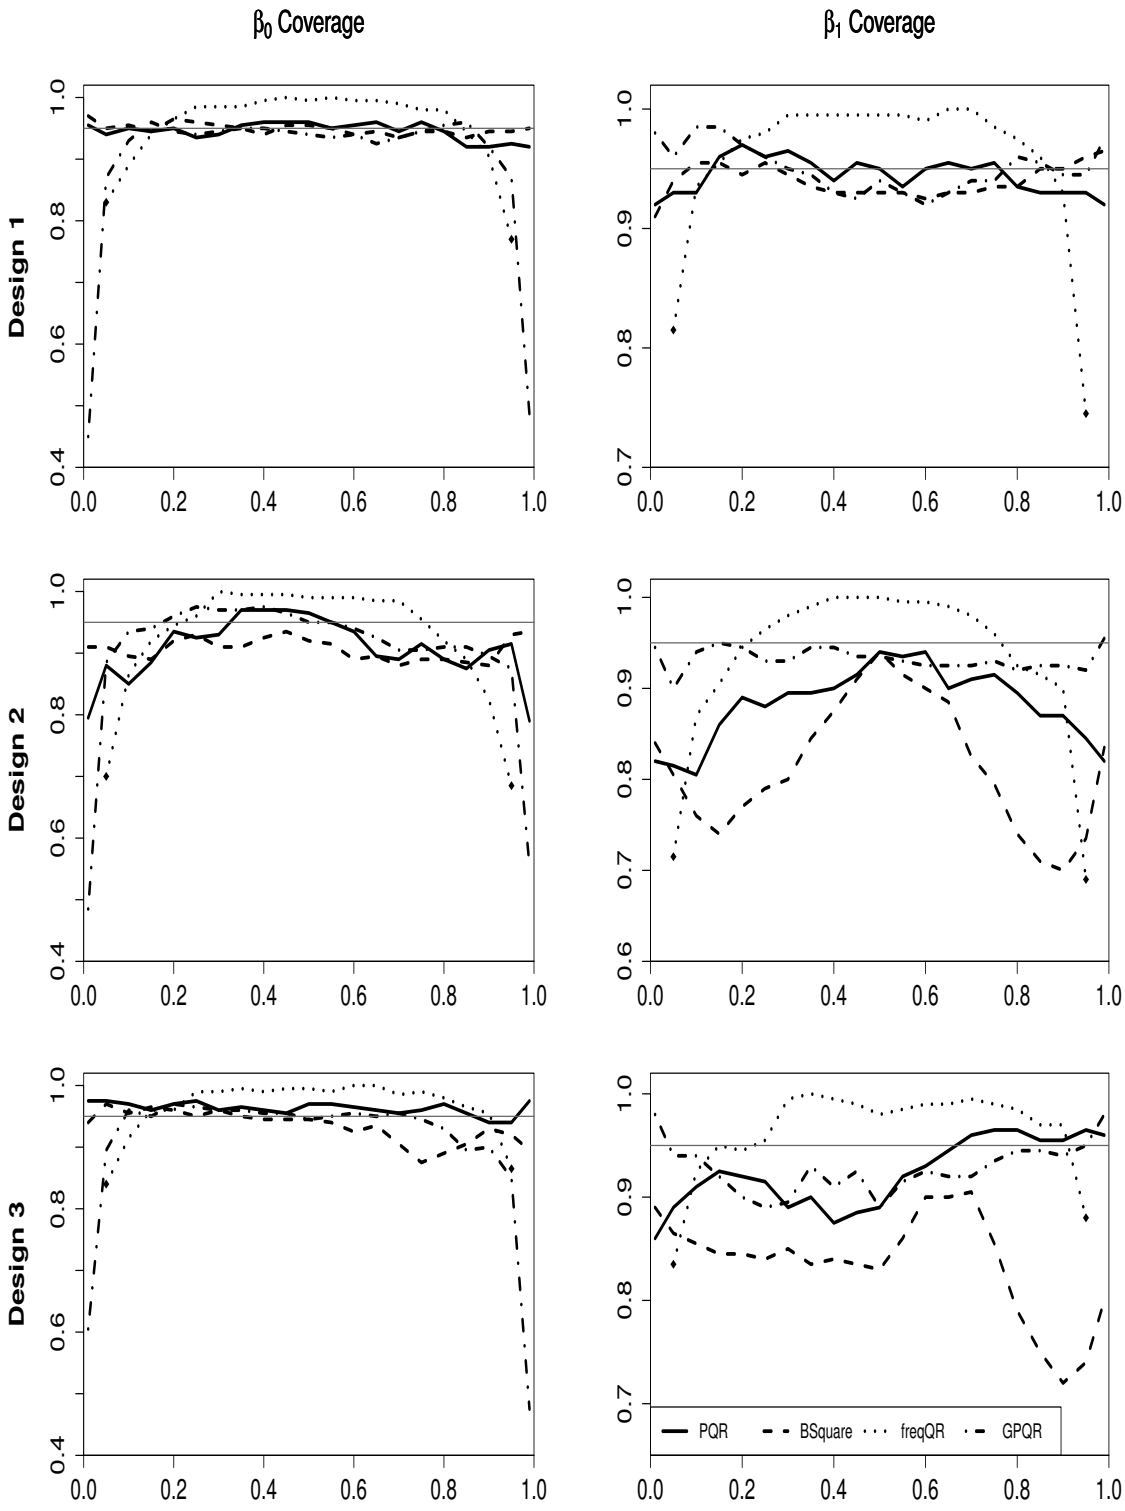

Figure 5: 95% Coverage probabilities for parameters β_0 (left) and β_1 (right) at $\tau = 0.01, 0.05, 0.1, \ldots, 0.95, 0.99$. freqQR coverage probabilities at $\tau = 0.01$ and 0.99 are not available, so this curve is truncated at $\tau = 0.05$ and 0.95 (diamond points).

GPQR has poor coverage probabilities for β_0 for extreme quantiles, for which the method also presented high RMSE (Figures 4 and 5). Prior complexity naturally compromises model interpretability and usage, which is a disadvantage of this approach. Prior information might be affecting estimation here, although default settings were used. From Figure 5, we can also see that freqQR coverages are generally too wide for middle quantiles and too narrow at the extremes (for $\tau = 0.05$ and 0.95, as the more extremes are not available). The BSquare approach performed poorly for some of the parameters. PQR has, in general, nice coverage probabilities compared to the alternative approaches.

For the multivariate design 4, we considered the estimation at quantile levels $\tau = 0.01, 0.05, 0.50$ with $N = 350$ samples. PQR was fitted based on 150.000 MCMC draws and burn-in of 50.000. For the other methods, previous configurations were adopted. RMSE and coverage results are presented in Tables 1 and 2, respectively.

Table 1: RMSE ($\times 100$) for design 4

	β_0	β_1	β_2	β_3	β_4	β_5
$\tau = 0.50$						
PQR	13.00	24.74	21.15	20.74	20.29	19.19
BSquare	16.13	28.20	26.52	19.85	18.91	19.00
freqQR	13.64	22.37	23.47	21.71	21.14	22.05
GPQR	13.37	24.86	23.00	22.61	21.84	21.03
$\tau = 0.05$						
PQR	21.31	30.61	35.89	32.16	34.63	31.55
BSquare	22.62	65.06	83.90	19.78	19.51	19.32
freqQR	21.86	37.57	39.69	35.90	38.68	37.01
GPQR	20.85	32.44	36.93	32.62	34.58	30.18
$\tau = 0.01$						
PQR	32.73	40.03	48.45	39.74	43.11	38.49
BSquare	65.67	72.44	91.51	19.77	19.50	19.31
freqQR	39.31	52.41	57.09	52.83	54.29	51.39
GPQR	45.67	50.27	57.16	48.12	50.34	46.58

BSquare had issues in estimating the parameters for this multivariate design. In particular β_0, β_1 and β_2 presented high RMSE's and low coverages, as shown in Tables 1 and

Table 2: 95% Coverage probabilities for design 4

	β_0	β_1	β_2	β_3	β_4	β_5
$\tau = 0.50$						
PQR	0.94	0.90	0.95	0.95	0.96	0.96
BSquare	0.88	0.82	0.90	0.96	0.94	0.94
freqQR	0.98	1.00	0.99	0.98	1.00	0.98
GPQR	0.96	0.86	0.89	0.92	0.95	0.92
$\tau = 0.05$						
PQR	0.96	0.98	0.94	0.98	0.96	0.96
BSquare	0.96	0.19	0.09	0.96	0.94	0.94
freqQR	0.90	0.88	0.86	0.89	0.85	0.90
GPQR	0.92	0.92	0.92	0.96	0.92	0.94
$\tau = 0.01$						
PQR	0.97	0.96	0.91	0.96	0.96	0.98
BSquare	0.72	0.14	0.08	0.96	0.94	0.94
freqQR	0.41	0.52	0.46	0.47	0.48	0.50
GPQR	0.76	0.96	0.97	0.97	0.98	0.97

2. In fact, the estimated quantile planes corresponding to different quantile levels were generally parallel, which obviously impacted the estimation of all parameters that vary with τ. This drawback of the non-crossing constraints imposed in Reich and Smith (2013) often happens for multivariate examples, unless large samples are available so that crossing occurs infrequently.

From Table 1, PQR has generally the smallest RMSE, significantly outperforming GPQR at $\tau = 0.01$ and also notably better than freqQR for $\tau = 0.05$ and $\tau = 0.01$. Moreover, among all methods, PQR has coverages closest to the nominal level. As noted before, freqQR has coverages consistently above the nominal level at $\tau = 0.50$ and below it at the extremes ($\tau = 0.05$ and $\tau = 0.01$). Again GPQR has poor coverage for β_0 for extreme quantiles.

Note that PQR has great performance despite the small number of pyramid levels ($M = 2$, $\tau = 0.01, 0.05, 0.50$). Indeed, increasing M does not significantly affect the results, corroborating the proximity between the least false and true parameter values.

5 Real examples

In this section, we illustrate the proposed method on two real examples involving quantile linear splines and extremal quantile modelling.

5.1 Quantile linear splines

The PQR model can be easily extended to the scope of piecewise linear quantile regression, by placing one pyramid at each spline knot. In order to illustrate this extension, we consider the famous heteroscedastic dataset of light detection and ranging (lidar) described in (Holst et al. 1996). This dataset is used for measuring the concentration of mercury in the atmosphere, and contains $N = 221$ observations. As concentration depends on altitude we use the distance range, *i.e.* the distance travelled before a laser light is reflected back to its source, as covariate. Since the concentration is inversely proportional to the first derivative of the dependent variable logratio (the logarithm of the ratio of received light from two laser sources, see Ruppert et al. 1997 for more details), piecewise linear fitting is appealing.

Here estimation of quartiles $\tau = 0.25, 0.50, 0.75$ is considered using 7 equally spaced knots. The piecewise linear quantile curves were estimated using 60.000 MCMC samples and 10.000 burn-in. PQR linear splines fitting is illustrated in Figure 6.

Without relying on strong parametric assumptions about the form of conditional distributions, PQR is able to simultaneously and flexibly fit lidar quartiles, as shown in Figure 6. The linear fitting provides point and interval estimates of mercury concentration quartiles (given the slope estimates) for different atmospheric altitudes, providing precious information for monitoring atmospheric pollutants.

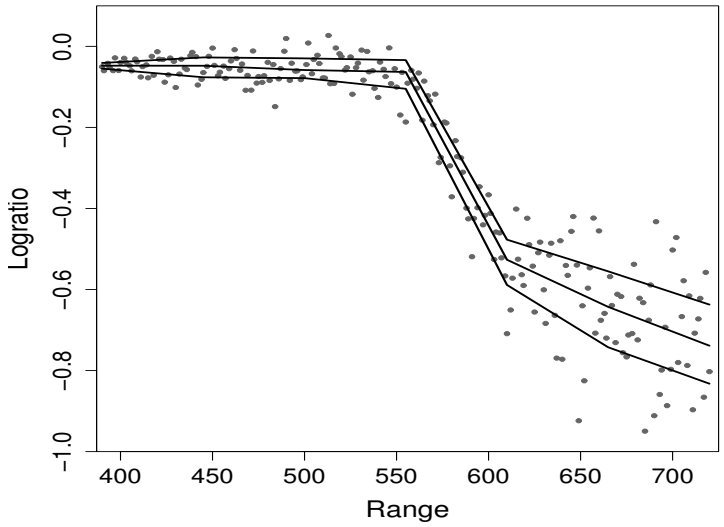

Figure 6: Linear splines fitting from PQR for lidar example at $\tau = 0.25, 0.50, 0.75$.

5.2 Extreme quantile modelling

In extreme value analysis, it is common practice to use the so-called extreme value distributions to make inference on the tails of the distribution of the data. Using a parametric model places strong assumptions on the data, but is an attractive approach since data is often scarce in the extremal regions. However, a long standing issue is the fidelity of the data to the parametric assumptions, see Coles (2001). We propose in this application to model linear quantiles of extreme data using PQR, that allows us to drop these parametric assumptions, but instead use the information from extreme distributions as prior when centring the quantile process.

Here we will apply PQR to model extreme tropical cyclones. The dataset consists of 82 observations of cyclones whose wind speed is greater than 96 knots (kt) threshold, recorded in the US coast from 1899 to 2006 (this is an updated version of the data analysed in Jagger and Elsner 2009 which included 79 cyclones; the update is available in the authors' webpage). Jagger and Elsner (2009) considers that these data follow a Generalized

Pareto Distribution (GPD), with cdf given by

$$G(y) = 1 - [1 + \xi(y - \mu)/\sigma]_+^{-1/\xi},$$

where $(h)_+ = max(h, 0)$, μ is the fixed threshold, and $\sigma > 0$ and ξ are the scale and shape parameters respectively.

Therefore, here we consider fitting PQR using GPD as the quantile process centring distribution. Similarly to the gaussian case we assume here that the unknown parameters (σ, ξ) change linearly in x. Following Jagger and Elsner (2009), we model extreme tropical cyclone wind speed quantiles for quantile levels $\tau = 0.10, 0.25, 0.50, 0.75, 0.90$ as a function of the Southern Oscillation Index (SOI) and the sunspot number (SSN), both averaged over August-October and standardised. We used for the estimation 60.000 MCMC draws and burn-in of 10.000. Figure 7 presents the parameter estimates and 90% confidence interval for PQR, obtained as the upper and lower 0.05 sample quantiles of the posterior samples. For comparison, freqQR estimates are also indicated (dashed lines).

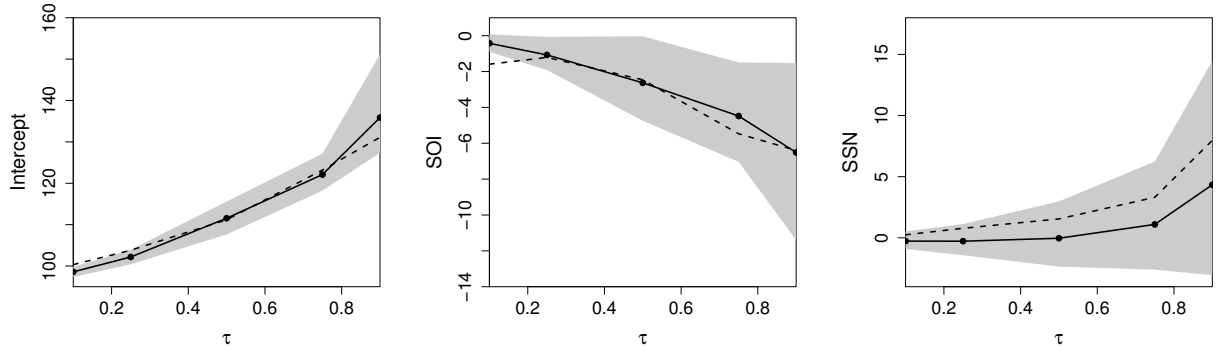

Figure 7: Parameter estimates for extreme tropical cyclone example using PQR (solid line) and freqQR (dashed line) at $\tau = 0.10, 0.25, 0.50, 0.75, 0.90$. The grey shading indicates 90% confidence interval for PQR.

As illustrated in Figure 7, the wind speed increases with decreasing SOI, which is expected as small SOI is associated with El Nino warming events, which in turn favour extreme cyclones, as explained in Jagger and Elsner (2009). From SOI parameter estimates'

plot, we can also see that PQR provides smoother and nicer estimates than freqQR, which lack borrowing strenght from the neighbours τ. As in Jagger and Elsner (2009), SSN is generally positive associated with extreme winds, but this is not a statistically significant association.

Modelling extreme winds is paramount for predicting the chance of extreme cyclones under different climate conditions, providing valuable source of information for decision makers. The proposed nonparametric approach for inference in extreme quantile modelling is appealing, as it is flexible, robust and provides meaningful parameter estimates.

6 Discussion

This paper extends the quantile pyramids introduced by Hjort and Walker (2009) to the regression context, proposing a novel simultaneous linear quantile regression model, named pyramid quantile regression (PQR).

PQR avoids strong parametric assumptions about the conditional distributions, which adds great modelling flexibility and circumvents the need to make parametric assumptions about the distribution of the data. In addition, the model is parametrised in terms of the quantiles themselves, this is a natural way of modelling quantile regression and allows for easy interpretation and incorporation of prior information. For instance, one can centre the conditional quantile priors on chosen distributions based on prior knowledge. We considered centring it on the Normal distribution, and showed that this choice by default works well for a variety of cases, including mildly asymmetric densities. Additionally, PQR can be used for flexible extreme quantile modelling by centring the prior on an extreme distribution, as opposed to strictly requiring the data to follow the parametric assumption, as is often the case in extreme value modelling. We illustrated this application in the modelling of extreme tropical cyclone winds in the US coast using pyramid prior centred on the Generalised Pareto Distribution (GPD). PQR can also be applied to linear splines fitting, as illustrated here through the famous lidar dataset.

Moreover, for fixed pyramid level M, PQR has posterior consistency to the least false parameters. In this context, simulation studies demonstrate that the proposed approach has small error and good coverage, overall outperforming its competitors, especially for extreme quantiles. In addition, consistency to the truth is achieved by considering growing pyramid level M with sample size, and our experience with empirical studies shows that M does not need to be large to obtain reasonable results, suggesting that the least false estimators are close to the truth.

In summary, PQR is a novel technique for fitting single or multiple quantiles based on a flexible nonparametric model, providing robust estimates with small errors and great coverages properties.

Acknowledgements

TR is funded by CAPES Foundation via the Science Without Borders (BEX 0979/13-9).

References

Bondell, H. D., B. J. Reich, and H. Wang (2010). Noncrossing quantile regression curve estimation. *Biometrika 97*(4), 825–838.

Chernozhukov, V., I. Fernandez-Val, and A. Galichon (2009). Improving point and interval estimators of monotone functions by rearrangement. *Biometrika 96*, 559–575.

Coles, S. G. (2001). *An introduction to statistical modeling of extreme values*. Springer Verlag, London.

Dette, H. and S. Volgushev (2008). Non-crossing non-parametric estimates of quantile curves. *Journal of Royal Statistical Society B 70*, 609–627.

Fang, Y., Y. Chen, and X. He (2015). Bayesian quantile regression with approximate likelihood. *Bernoulli 21*(2), 832–580.

Ferguson, T. S. (1974). Prior distributions on spaces of probability measures. *Annals of Statistics 2*, 615–629.

Hall, P., R. C. L. Wolff, and Q. Yao (1999). Methods for estimating a conditional distribution function. *Journal of American Statistical Association 94*, 154– 163.

Hanson, T. and W. O. Johnson (2002). Modelling regression error with a mixture of Pólya trees. *Journal of American Statistical Association 97*(460), 1020–1033.

He, X. (1997). Quantile curves without crossing. *American Statistician 51*, 186–192.

Hjort, N. L. (1992). On inference in parametric survival data models. *International Statistical Review 60*(3), 355–387.

Hjort, N. L. and S. G. Walker (2009). Quantile pyramids for Bayesian nonparametrics. *Annals of Statistics 37*(1), 105–131.

Holst, U., O. Hossjer, C. Bjorklund, P. Ragnarson, and H. Edner (1996). Locally weighted least squares kernel regression and statistical evaluation of lidar measurements. *Environmetrics 7*, 401–416.

Jagger, T. H. and J. B. Elsner (2009). Modeling tropical cyclone intensity with quantile regression. *International Journal of Climatology 29*, 1351–1361.

Koenker, R. (2005). *Quantile regression*, Volume 38 of *Econometric Society Monographs*. Cambridge: Cambridge University Press.

Koenker, R. and J. Bassett, Gilbert (1978). Regression quantiles. *Econometrica 46*(1), 33–50.

Kottas, A. and A. E. Gelfand (2001). Bayesian semiparametrc median regression modelling. *Journal of American Statistical Association 96*, 1458–1468.

Kottas, A. and M. Krnjajić (2009). Bayesian semiparametric modelling in quantile regression. *Scandinavian Journal of Statistics 36*, 297–319.

Lavine, M. (1992). Some aspects of Pólya tree distributions for statistical modelling. *Annals of Statistics 20*(3), 1222–1235.

Lavine, M. (1994). More aspects of Pólya tree distributions for statistical modelling. *Annals of Statistics 22*, 1161–1176.

Moss, J. (2015). Cube root asymptotics. Master's thesis, Department of Mathematics, University of Oslo.

R Core Team (2014). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Reich, B. J., H. D. Bondell, and H. J. Wang (2008). Flexible Bayesian quantile regression for independent and clustered data. *Biostatistics 11*, 337–352.

Reich, B. J., M. Fuentes, and D. B. Dunson (2011). Bayesian spatial quantile regression. *Journal of the American Statistical Association 106*(493), 6–20.

Reich, B. J. and L. B. Smith (2013). Bayesian quantile regression for censored data. *Biometrics 69*, 651–660.

Rodrigues, T. and Y. Fan (2016). Regression adjustment for noncrossing Bayesian quantile regression. *Journal of Computational and Graphical Statistics*. (in press).

Ruppert, D., M. P. Wand, U. Holst, and O. Hössjer (1997). Local polynomial variance function estimation. *Technometrics 39*(3), 262–273.

Smith, L. and B. Reich (2013). *BSquare: Bayesian Simultaneous Quantile Regression*. R package version 1.1.

Sriram, K., R. V. Ramamoorthi, and P. Ghosh (2013). Posterior consistency of bayesian quantile regression based on the misspeficied asymmetric Laplace density. *Bayesian Analysis 8*(2), 1–26.

Strawderman, R. L. and A. A. Tsiatis (1996). On consistency in parameter spaces of expanding dimension: An application of the inverse function theorem. *Statistica*

Sinica 6, 917–923.

Tokdar, S. (2015). *qrjoint: Joint Estimation in Linear Quantile Regression*. R package version 0.1-1.

Tokdar, S. T. and J. B. Kadane (2012). Simultaneous linear quantile regression: a semiparametric Bayesian approach. *Bayesian Analysis 7*(1), 51–72.

Yang, Y. and S. Tokdar (2016). Joint estimation of quantile planes over arbitrary predictor spaces. *Journal of the American Statistical Association*. (in press).

Yu, K. and R. A. Moyeed (2001). Bayesian quantile regression. *Statist. Probab. Lett. 54*(4), 437–447.